Jl. of Educational Multimedia and Hypermedia (2004) **13**(4), 391-404

Opportunities for New "Smart" Learning Environments Enabled by Next-Generation Web Capabilities

PHILIP DODDS
Randall House Associates, Inc., USA
pdodds@rhassociates.com

J. D. FLETCHER
Institute for Defense Analyses, USA
dfletche@ida.org

Empirical evaluations suggest that use of interactive technologies can reduce the costs of instruction by about one-third. In addition, they can either increase achievement by about one-third while holding time constant or reduce time needed to achieve targeted instructional objectives by about one-third. These technologies can be delivered over the Web, which can also support systems that generate instruction on demand. Development of either generative instruction or pre-specified interactions will benefit from a ready supply of instructional objects such as those specified by the Sharable Content Object Reference Model (SCORM), which is now receiving wide, international acceptance. SCORM will be further enhanced by the development of the Semantic Web, which will allow more extensive links between available representations of knowledge and enhance the discovery of learning objects for use in instruction.

How might the emerging standards and infrastructure for Web-based learning objects enable, or even encourage, the development of sophisticated learning environments that evolve into *intelligent* tutoring systems? This question, which is the basic topic of our paper, raises many questions of its own. Suppose these learning objects do enable development of intelligent tutoring systems. Would the result be worthwhile? What do we mean by intelligent tutoring systems? What emerging standards, infrastructure, and learning objects are needed for these systems? How might the Web and Web services influence this evolution? And where does this line of development lead those of us who are concerned with human learning and capabilities? We try to address each of these questions as well (and as briefly) as we can in the following comments.

DISCUSSION

Are Technology-Based Learning Environments Worthwhile?

As discussed in more detail by Fletcher (2003), the case for using technology to create these learning environments may be summarized as the following:

(1) Tailoring instruction (education and training) to the needs of individual students has been found to be an instructional imperative and an economic impossibility. Research has determined that students tutored one-on-one score about 2 standard deviations higher on end-of-course achievement tests than students taught in one-on-many classrooms (Bloom, 1984). However, except for a few critical skills (e.g., airplane piloting, surgery), we cannot afford the one instructor for each student that such tutoring requires.

(2) In many situations, technology-based instruction can make this instructional imperative affordable. Under any appreciable student load, it is less expensive to provide instruction with technology than to hire a tutor for each student.

(3) Technology-based instruction has been found to be more effective than current classroom instructional approaches in many settings across many subject matters. Review of 233 evaluations of typical technology-based instruction found an average improvement of 0.39 standard deviations over classroom instruction, which is roughly equivalent to raising the performance of 50^{th} percentile students to that of 65^{th} percentile students. Review

of 44 similar evaluations of interactive multimedia instruction found an average improvement of 0.50 standard deviations, roughly raising the performance of 50th percentile students to the 69th percentile. Intelligent tutoring systems have produced improvements of 1.05 standard deviations, roughly raising 50th percentile performance of students to about the 85th percentile.

(4) Technology-based instruction is generally less costly than current instructional approaches, especially when many students or expensive devices are involved. In 16 studies where achievement under technology-based training was at least equal (and mostly superior) to that of classroom instruction, the cost ratios of the former to the latter were found to be 0.43 for initial investment, 0.16 for operating and support, and 0.35 overall.

(5) Technology-based instruction has been found to decrease the time needed to reach targeted instructional objectives. A review of 40 studies found that savings in the time needed to achieve given instructional objectives averaged about 30 percent. Hundreds of millions of dollars would be saved in Department of Defense specialized skill training if training time could be reduced by 30 percent. Time savings may be even more important in K-16 education where opportunities for students to expand their capabilities and develop their potential as rapidly as possible lays the foundation for fully-realized, satisfying lives as well as global competitiveness and economic health.

(6) Technology-based instruction has been found to be a cost-effective alternative for achieving instructional goals. Compared to reducing class size, providing professional tutors, using peer tutors, or increasing the length of the school day, the costs to provide technology-based instruction for 15 minutes each day were found to be the least expensive means for raising comprehensive mathematics scores on standardized tests.

(7) Technology-based instruction will become increasingly affordable and instructionally effective with the development and use of standardized, reusable, instructional objects. Early results have already indicated significant savings (Dodds, 2002).

(8) The knowledge structures underlying technology-based instruction can be readily (i.e., inexpensively) used to provide interactive performance aids that both lower training costs and enhance job performance (Fletcher & Johnston, 2002).

Overall, a *rule of thirds* emerges from assessments of technology-based instruction. Use of these technologies reduces the cost of instruction by about one-third. In addition, either it reduces time required for instruction by about one-third or it increases the amount of skills and knowledge acquired by about one-third.

It should be emphasized that technology-based instruction can be used either by individuals or by groups of individuals working in collaboration. It can be used in residential classrooms, remote classrooms, or any remote (distributed) location—workplace, home, or elsewhere. It can be available anytime, anywhere.

What Do We Mean By Intelligent Tutoring Systems and What Do They Add?

The features that garden-variety technology-based instruction provides are notable. It can: (a) accommodate an individual student's rate of progress, allowing as much or as little time as the student needs to reach instructional objectives; (b) tailor both the content and the sequence of instructional content to each student's needs; (c) make the instruction as easy or difficult, specific or abstract, applied or theoretical as necessary; and (d) adjust to students' most efficient learning styles (collaborative or individual, verbal or visual, etc.). These capabilities have been available and used in technology-based instruction from its inception in the 1950s (Fletcher & Rockway, 1986).

Intelligent in an intelligent tutoring system refers as much to intentions as to results. However, it is more than a marketing term. It refers to specific capabilities that have been the goals of intelligent tutoring systems development since such development was first attempted in the 1960s (Carbonell, 1970; Brown, Burton, & DeKleer, 1982). Two defining capabilities of intelligent tutoring systems are that they:

- Allow either the system or the student to ask open-ended questions and initiate instructional and mixed-initiative dialogue as needed or desired.
- Generate instructional material and interactions on demand rather than require developers to foresee and pre-store all such materials and interactions needed to meet all possible eventualities.

Mixed-initiative dialogue requires a language for information retrieval, decision-aiding, and instruction that is shared by both the system and the student/user. Natural language has been a frequent choice for this capability (e.g., Brown, Burton, & DeKleer, 1982; Graesser, Person, & Magliano, 1995), but the language of mathematics, mathematical logic, electronics, and other well-structured communication systems have been used (Suppes, 1981; Sleeman & Brown, 1982; Psotka, Massey, & Mutter, 1988).

The generative capability requires the system to devise on-demand interactions with students—not draw from predicted and pre-stored formats. This capability involves not just generating problems tailored to each student's needs, but also coaching, hints, critiques of completed solutions, appropriate and effective teaching strategies, and, overall, the interactions and presentations needed for one-on-one tutorial instruction. These interactions must be generated from information primitives using an *instructional grammar* that is analogous to the deep structure grammar of linguistics.

Motivations for these two capabilities can be found in basic research on human learning, memory, perception, and cognition. Findings from this research have led us to view all cognitive processes as constructive and regenerative (Neisser, 1967). They have extended general theories of perception and learning from the fairly-strict, logical positivism of behavioral psychology, which emphasized directly observable actions, to consideration of the internal, cognitive processes that are assumed to mediate and enable human learning. The hallmark of these conceptions of cognition is that seeing, hearing, and remembering are all acts of *construction,* making more or less use of the limited stimulus information provided by our perceptual capabilities.

The generative capability sought by intelligent tutoring systems developers is not something merely nice to have, but essential if we are to advance beyond the constraints of prescribed, pre-branched, programmed learning, and the ad-hoc principles commonly used now to design technology-based instruction. We need it if we are to deal successfully with the immensity, extent, and variability of human cognition.

Are Standards, Infrastructure, and Learning Objects Needed?

Specification for learning objects has become an essential component of the Advanced Distributed Learning (ADL) Initiative. This initiative is the most recent and visible effort in a long campaign to adopt the benefits of technology-based instruction and performance-aiding in routine practice. It is intended to accelerate large-scale development of dynamic and cost-effective learning software and stimulate a vigorous market for learning software. Its goal is to ensure access to high quality learning (education, training, and performance-aiding), tailored to individual needs and capabilities, and available at any time and any place (Dodds, 2002).

The ADL initiative is preparing for a future in which communication networks and personal delivery devices are pervasive, inexpensive, and effectively transparent to users through ease of use, expanded bandwidth, and portability. It will establish knowledge libraries, or repositories, where learning *objects* may be accumulated and cataloged for broad distribution and use. Because of their enhanced accessibility, these objects will be ready for assembly on-demand and in real time into instructional and performance-aiding materials that are tailored to the capabilities, intentions, and learning state of each individual or group of individuals who require them (Dodds, 2002). ADL and intelligent tutoring systems, therefore, have a number of key goals in common:

- Both are generative in that they seek to prepare and present interactions on-demand and in real time.
- Both are intended to provide instructional interactions that are tailored in content, sequence, difficulty, style, etc. to users' intentions, backgrounds, and needs.
- Both have a stake in research intended to accomplish such individualization.
- Both can be used equally well in instruction and performance-aiding.
- Both are intended to accommodate mixed-initiative dialogue in which either the technology or the user initiates or responds to open-ended inquiry and discussion.
- Both require a supply of sharable instructional objects readily available for the generation of instruction or performance aiding presentations.

To date most of the ADL effort has been devoted to the specification of instructional objects that will populate learning libraries and other Web-available repositories. These objects are separated from context-specific run-time constraints and proprietary systems so that they can be incorporated into other applications. They have common interfaces and data exchange formats. They are *accessible* so that they can be indexed and readily found or *discovered, interoperable* so that they operate across a wide variety of hardware, operating systems, and Web browsers, *durable* so that they do not require modification as versions of the underlying software systems change, and *reusable* so that they can be adapted and used by many different development tools. The ADL initiative has coordinated groups of industry, academic, and government stakeholders working together to specify objects that meet these criteria. These specifications have produced evolving, cumulative versions of SCORM, the Sharable Content Object Reference Model (Dodds, 2002).

As presently defined, SCORM objects can be entire courses, lessons within courses, or modules within lessons—their granularity or size remains at issue. ADL development is presently focused on packaging these objects for what has been called the *educational object economy* (Spohrer, Sumner, & Shum, 1998). One idea behind such an economy is that the emphasis in preparing materials for technology-based instruction (or decision-aiding) will shift from the current concern with preparing content components or instructional objects to one of integrating already available content into meaningful and relevant presentations.

Many technicians, software engineers, instruction designers, and cognitive researchers who come from all sectors of the economy in the Americas, Europe, and Asia have joined in this quest. The task of specifying and developing these objects has become a global effort. The primary contribution of ADL has been to orchestrate this effort and document its results.

Towards More Adaptive and Intelligent Learning Systems

Until recently, most mainstream learning systems have relied on predetermined and often fixed-path delivery of content. Such systems lack agility in adapting to learners' mastery states, and are thereby limited in their ability to tailor learning experiences to individual learners. An adaptive, intelligent learning system needs an accurate model of the learner, a model of the knowledge domain, and a capability that can evaluate the differences between the two models. It must be able to identify or devise instructional strategies that will achieve desired instructional outcomes on-demand and in real time.

SCORM presently provides a rules-based *learning strategy* that enables Sharable Content Objects (SCOs) to set the state of globally accessible records. These records can store the learner's degree of mastery in the form of a score or a pass/fail state, or they may store the progress of the learner in terms of completion. A *hook* was included in the records permitting them to reference externally defined competencies. As the learner is sequenced through the SCOs, the learning system builds up a representation of the learner's mastery and progress. The objective records may be viewed as a simple model of the learner's state (Gibbons & Fairweather, 1998). How far SCORM's sequencing, navigation, and assessment capabilities can be pushed in the direction of generative intelligent learning systems remains an empirical issue currently being settled by the development of increasingly sophisticated learning systems that observe SCORM specifications.

Another emerging specification called *IMS Reusable Definition of Competency or Educational Objective* (2002) defines a means of building a taxonomy of competency definitions that meet specific objectives. This taxonomy may be organized hierarchically to represent dependencies, supporting skills, or prerequisites. Each competency definition has a text description of the competency and a unique identifier that may be referenced externally. The organization of a competency definition could represent specific skills or knowledge to be acquired for a specific task or subject domain (e.g., as one might find in Quantitative Domain Mapping). Since objectives records in SCORM can reference the competency model identifiers, the means to compare the state of the learner and the desired competencies now exists. This capability provides a system-based means to perform knowledge and skills gap analyses leading to more sophisticated and adaptive strategies that use such information (Wiley, 2000).

As learning system specifications become more robust, they will also become more adaptive. Improved assessment methods and results are emerging that will continuously and unobtrusively extract information from instructional interactions and better represent the state of the learner (e.g., Conati, Gertner, VanLehn, & Druzdel, 1997; Corbett, Koedinger, & Anderson, 1997). The strategies developed by learning systems will be further enhanced by learner profile information, which can *pre-load* the learner model with mastery information from outside of the system. This process will enhance the processes used by technology-based instructional systems to bypass relevant content of pre-mastered material (e.g., holders of certificates in particular subjects) and concentrate on relevant material yet to be learned.

The emerging specifications have enabled a means of modeling and tracking the learner and referencing that model's external competency/knowledge models. The specifications now allow conditional rules that can tailor what the learner experiences to his/her mastery and progress (e.g., Dodds, 2002). Future services and processes will extend these basic capabilities in more sophisticated and nuanced ways.

Basically, what we seek is an engineering of instruction in which outcomes such as retention of skills and knowledge, transfer of learning to new but similar applications, motivation to continue study, speed of response, accuracy of response, and so forth are achieved reliably by all learners. Such engineering would adjust and modulate the learning experience for each individual. It would identify and devise learning strategies as described earlier, but then, on-demand and in real time, locate objects for each successive interaction with the learner that are appropriate to the outcome being sought (e.g., the learner's characteristics, level of knowledge, and style of learning, and the instructional strategy that was identified or devised). This is a significant challenge for both instructional objects and Web-based services, but current work suggests that they may successfully meet it.

How Might the Web and Web Services Influence This Evolution?

One way the current and near term capabilities of learning systems might evolve is through the Semantic Web, which will provide powerful new technologies for both knowledge representation and the ontologies needed to connect them (Berners-Lee, Hendler, & Lassila, 2001). These technologies will provide ways not only to relate but also to reason about information from widely different domains.

The Semantic Web is intended to imbue information available on the Web with sufficient meaning to substantially improve the cooperation between computers and human beings. It requires abstract representation of information on the Web using a Resource Description Framework and other specifications yet to be developed. Dealing with the semantic content of Web pages and information will enhance the process of discovery needed to access relevant information and objects from the Web. Through an ontology, consisting of a taxonomy and a set of inference rules that formally define operations and relations among terms, it will be possible to identify and expose semantic linkages between highly disparate bodies of information.

If successful, the Semantic Web will integrate real-world knowledge and skills acquired through simulation, education, training, performance-aiding, and experience. It will provide a foundation for building more comprehensive and substantive models of subject matter domains and learners' levels of mastery than we now have and combine them with more precise discovery of the instructional objects needed to produce desired human competencies. Learners and practitioners will be presented with a constellation of related activities—learning, doing, trying, and referencing. This integration, combined with the already available functionalities of intelligent tutoring systems, provides the basis for a next-generation meta-architecture and learning environments based on instructional objects.

Core components of the Semantic Web will be built on top of existing and emerging Web standards. These standards provide the means to express complex relationships and inference rules that are processed by Web services to perform specific tasks such as profiling learners, representing their skills, knowledge, and abilities, linking these representations to instructional objects, and managing their progress toward objectives and competencies. Web services will serve as reusable, black-box applications that generate other Web-based applications from objects. They will use open Internet standards, such as Hyper-Text Transfer Protocol (HTTP), Extensible Markup Language (XML), Universal Description Discovery and Integration (UDDI), and Simple Object Access Protocol (SOAP), to exchange information between applications as needed. The services will be language, platform, and object model independent. They will enable different applications running on different operating systems, developed with different object models using different programming languages to cooperate and become easily-used Web applications. They will provide flexible, standards-based capabilities for binding applications together over the Internet, taking advantage of existing infrastructure and applications.

Content Object Discovery and Retrieval

Searching and discovering contextually relevant instructional content have become major topics. The success of Google and other Web search engines have whetted everyone's appetite for *just-when-you-want-it* search and retrieval and have demonstrated the value and utility of content discovery. Presently, Google may be the single most important, effective, and widely used source of Web-based education. However, Google's method of locating content, by text crawling and indexing only and by retrieving anything that is available and remotely relevant, has limited its use as a discovery system for just-in-time focused, content assembly. Its operation could be substantially improved if it were to *cooperate* with content and retrieve only what is intentionally prepared and published for discovery.

A series of Internet specifications that have been in development for some time appear to be maturing and might form a framework on which sophisticated search, authentication, accreditation, and resolution services might be built. These services will produce comprehensive means for instructional programs to locate appropriate content and then access it. There are at least two capabilities that are needed to build such services: (1) content object identification and resolution, and (2) discovery indexing with search criteria.

Content identification (to select candidate objects) and resolution (to narrow the selection to objects that are precisely relevant) are being addressed through the use of Universal Resource Names (URNs), which are intended to serve as persistent, location-independent, resource- identifiers. The Corporation for National Research Initiatives (CNRI) has created a URN implementation called The Handle System (Kahn & Wilensky, 1995; CNRI, 2003) that allows digital objects to obtain a unique identifier and link each object to its location—wherever that might be – through the use of a Handle Resolution Service (similar to how domain names resolve Internet protocol addresses through the Domain Name System). CNRI hosts a global root server that can be queried during resolution requests. The Handle System addresses a key repository problem: the unique identification of objects along with their present location and descriptive metadata.

Discovery indexing presents another challenge. How does one locate relevant content in the first place? One approach may be to use the Handle System to set up a registry and index of content repositories so that a repository may be located and searched. To do so, some form of external search means must be enabled against some defined search criteria. The Global Information Locator Service (GILS) specification, which is based on an International Standards Organization search standard (ISO 23950), addresses how repositories might identify (discover) content within a repository collection.

The Common Indexing Protocol (CIP), developed by the Internet Engineering Task Force (IETF), allows the owner of content to create its index metadata while also allowing this indexing information to be shared among different servers. This enables the development of new search and discovery services. There are new learning and performance-aiding specifications emerging that permit the identification of skills, competencies, and knowledge so that logical relations among them can be made and then represented in taxonomies that are relevant to specific but quite different communities of practices. One example of such specifications is the Reusable Definition of Competency or Educational Objective specification from the IMS Global Consortia, which is now being advanced at the Institute of Electrical and Electronic Engineers (IEEE) as a candidate standard.

These specifications show great promise for the design and deployment of Internet and Web services that will enable accurate, precisely-focused, and contextually-correct discovery and retrieval of learning content objects on a highly scalable basis—an ability still not yet available. They will allow instructional programs to continuously and unobtrusively assemble models of each learner's state of knowledge, style of learning, and progress toward instructional objectives. These models will in turn support the precise tailoring of instructional interactions to each student that is a characteristic and unique strength of one-on-one tutoring; they will provide an Aristotle for every Alexander and a Mark Hopkins for the rest of us. The next several years are likely to see a great deal of emphasis on developing the specifications and services needed to make this possible.

Where Might These Capabilities Take Us?

The emphasis on instructional technology brings us to revolutions in instruction. The first of these may have occurred with the development of written language about 7000 years ago. It allowed the content of advanced ideas and teaching to transcend time and place. The second revolution in instruction began with the technology of books. Books made the content of high quality instruction available anywhere and anytime, but also made such content inexpensive and thereby accessible to many more people. A third revolution in instruction appears to be accompanying the introduction of computer technology. The capability of this technology for real time adjustment of instructional content, sequence, scope, difficulty, and style to meet the needs of individuals suggests a third pervasive and significant revolution in instruction. It makes both the content *and* the interactions of high-quality instruction widely and inexpensively accessible—again anytime, anywhere.

Building on this possibility, ADL, SCORM, intelligent tutoring, and the Semantic Web will provide a foundation for generative education, training, and performance-aiding available anytime, anywhere. These developments will capitalize on the growth of electronic commerce and the global information grid. They will build on this worldwide, almost irresistible activity, accelerate it, and apply it to a full spectrum of education, training, and performance-aiding needs.

References

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. Scientific American, 284, 34-43.

Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*, 4-16.

Brown, J. S., Burton, R. R., & DeKleer, J. (1982). Pedagogical, Natural Language and Knowledge Engineering in SOPHIE I, II, and III. In D. Sleeman and J. S. Brown (Eds.). *Intelligent Tutoring Systems* (pp 227-282). New York, NY: Academic Press.

Carbonell, J. R. (1970). AI in CAI: An artificial intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems, 11*, 190-202.

CNRI The Handle System (2003). Available at: http://www.handle.net/introduction.html.

Conati, C., Gertner, A., VanLehn, K., & Druzdzel, M. (1997). On-line student modeling for coached problem solving using Bayesian Networks. In A. Jameson, C. Paris, & C. Tasso (Eds.) Proceedings of UM-97, Sixth International Conference on User Modeling. New York, NY: Springer Wien. Also available from: http://um.org.

Corbett, A. T., Koedinger, K. R., & Anderson, J. R. (1997). Intelligent tutoring systems. In Helander, M. G., Landauer, T. K., & Prabhu, P. V. (Eds.) *Handbook of Human-Computer Interaction*, (pp. 849-874). Amsterdam, The Netherlands: Elsevier Science.

Dodds, P. V. W. (Ed.) (2002). *Sharable Courseware Object Reference Model (SCORM) Version 1.2* (IDA Document D-2677). Alexandria, VA: Institute for Defense Analyses. Evolving versions of SCORM are available at: http://www.idanet.org.

Fletcher, J. D. (2003). Evidence for Learning from Technology-Assisted Instruction. In H. F. O'Neil Jr. & R. Perez (Eds.) *Technology Applications in Education: A Learning View* (pp 79-99). Hillsdale, NJ: Lawrence Erlbaum Associates.

Fletcher, J. D. & Johnston, R. (2002). Effectiveness and Cost Benefits of Computer-Based Aids for Maintenance Operations. *Computers in Human Behavior, 18*, 717-728.

Fletcher, J. D. & Rockway, M. R. (1986). Computer-based training in the military. In J. A. Ellis (Ed.) *Military Contributions to Instructional Technology* (pp 171-222). New York, NY: Praeger Publishers.

Gibbons, A. S. & Fairweather, P. G. (1998). *Computer Based Instruction: Design and Development*. Englewood Cliffs, NJ: Macmillan Library Reference.

Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology, 9*, 1–28.

IMS Reusable Definition of Competency or Educational Objective - Information Model: Version 1.0 Final Specification. IMS Global Learning Consortium, Inc. Available at: http://www.imsglobal.org/competencies/rdceov1p0/imsrdceo_infov1p0.html.

Kahn, R. & Wilensky, R. (1995). *A Framework for Distributed Digital Object Services*. Available at: http://www.cnri.reston.va.us/k-w.html.

Neisser, U. (1967). *Cognitive Psychology*. New York, NY: Appleton, Century, Crofts.

Psotka, J., Massey, L. D., & Mutter, S. A. (Eds.) (1988). *Intelligent tutoring systems: Lessons learned.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Sleeman, D., & Brown, J. S. (Eds.) (1982). *Intelligent tutoring systems.* New York, NY: Academic Press.

Spohrer, J., Sumner, T. & Shum, S. B. (1998). Educational authoring tools and the educational object economy: Introduction to the special issue from the East/West group. *Journal of Interactive Media in Education.* Available at: http://www-jime.open.ac.uk/98/10/spohrer-98-10-paper.html.

Suppes, P. (Ed.) (1981). *University-level computer assisted instruction at Stanford: 1968–1980.* Stanford, CA: Institute for Mathematical Studies in the Social Sciences.

Wiley, D. (2000). *The Instructional Use of Learning Objects.* Online at: http://www.reusability.org/read.